Oracle AI for Fusion Applications

How do I use AI Agent Studio?

Oracle AI for Fusion Applications
How do I use AI Agent Studio?

G36251-05

Author: Soniya Mathew

Contents

Get Help

There are a number of ways to learn more about your product and interact with Oracle and other users.

Get Help in the Applications

Some application pages have help icons to give you access to contextual help. If you don't see any help icons on your page, click your user image or name in the global header and select Show Help Icons. If the page has contextual help, help icons will appear.

Get Training

Increase your knowledge of Oracle Cloud by taking courses at *Oracle University*.

Join Our Community

Use *Cloud Customer Connect* to get information from industry experts at Oracle and in the partner community. You can join forums to connect with other customers, post questions, suggest *ideas* for product enhancements, and watch events.

Share Your Feedback

We welcome your feedback about Oracle Applications user assistance. If you need clarification, find an error, or just want to tell us what you found helpful, we'd like to hear from you.

You can email your feedback to *oracle_fusion_applications_help_ww_grp@oracle.com*.

Thanks for helping us improve our user assistance!

1 Overview

Overview of AI Agent Studio

AI Agent Studio is a design-time environment that empowers you to create, configure, validate, and deploy AI agents to meet your organization's needs.

With AI Agent Studio, you can easily extend preconfigured agent templates, and even build new agents and multiagent flows from scratch. AI Agent Studio is fully integrated into Oracle Fusion Cloud Applications, providing secure and seamless access to the knowledge stores, tools, and APIs of Fusion Applications. This integration enables agents to be deployed directly into the flow, ensuring an efficient process.

Key Capabilities of AI Agent Studio

Feature	Description
Agent template libraries	Use templates and natural language prompts to create or fine tune agents for common business scenarios, such as opportunity-to-quote processing and shift scheduling.
Agent team orchestration	Configure multiple agents to collaborate on multistep processes, integrating user approvals where necessary.
Agent extensibility	Change and extend existing agents in Oracle Fusion Cloud Applications by incorporating new data sources, prompts, and APIs to fit specific business needs or industry requirements.
Native integration with Fusion Applications	Directly access APIs and tools in Fusion Applications, ensuring seamless deployment of agents without complex modifications.
Third-party system integration	Connect with external systems and collaborate with third-party agents for end-to-end automation with secure API support.
Trust and security framework	Automatically applies the security configurations, policies, and access controls of Fusion Applications, ensuring compliance with enterprise security standards.
Validation and testing tools	Use built-in tools to make your agents reliable, repeatable, easy to explain, and secure by verifying AI-driven flows before deployment.

Components of AI Agent Studio

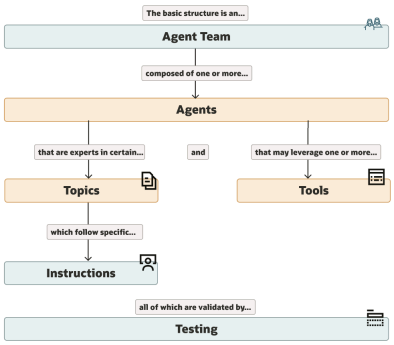

Component	What It Does	Example
Agent Team	Consists of a structured sequence of steps or actions that a single agent or a group of agents follows to accomplish a specific business task or answer a user query. The agent team is the component that can be deployed for use. • Consists of conversation logic, system integration, and user support flow. • Defines how the agent acts for a particular use case — what to do, when to do it, and how to respond based on user inputs or back-end system data.	A recruitment agent (that might consist of multiple agents) schedules interviews, screens resumes, calculates salary, and generates offers, based on policies and approvals.
Agents	Leverages a large language model to reason, create action plans, and interact with users to gather information and take direction. On behalf of users, the AI agent can do tasks that enhance productivity, efficiency, and the overall user experience. An agent must be added to an agent team, so that it can be deployed for use. AI agents can be categorized into various types. • User-proxy agent: Acts on behalf of a business user to provide input to another agent or group of agents. It's sometimes referred to as a conversational agent. • Supervisor agent: Orchestrates the use of agents within an agentic flow. • Specialist or utility agent: Focuses on a specific role or expertise and can be skilled in using a particular tool. AI agents can also have one or more of these characteristics: • Persona-based agent: Represents a specific role, such as benefits administrator, customer service representative, and finance administrator. • Tool user: Uses technology-related tools, such as calculators, web search queries, document embedding, and calendar schedulers. • Task-oriented agent: Understands their assignment or task, as a single agent or as part of a multiagent flow.	A scheduling or calendar agent that manages your workday by following your instructions. It can accept new calendar invites and propose alternative times when needed.

Component	What It Does	Example
Topics	Defines the areas of expertise through instructions that set the boundaries and constraints for agent conversations and abilities.	An employee benefits agent can contain topics such as Health Savings Account (HSA), retirement benefits, and stock plans.
Tools	Defines the additional utilities an agent can use to accomplish a task. One or more tools are assigned to agents, and they're reusable among agents.	• Calculator tool • Email tool • Business object tool • User query tool • Document retrieval tool for retrieval-augmented generation (RAG)
Instructions	Specifies the natural language rules that define the rules or conditions applied to a given topic. Instructions are part of the prompts that are sent to the underlying large language model. They can also contain guidelines and guardrails that set the parameters of an agent response.	Instructions for the payroll deduction topic: Make sure you've information regarding the number of dependents either by asking the user or querying the system. If you don't know the answer, don't make up a response.
Testing	Enables administrators to test the agent team design, ensure correct tool, topic, and instruction configurations, and validate reasoning and sources cited by the agent.	Provide example responses to a series of questions a user would likely ask an agent and details about how the agent arrived at its response.

2 Access Requirements for AI Agent Studio

Access Requirements for AI Agent Studio

You can give access to AI Agent Studio by assigning predefined duty roles to job roles. Also, make sure to complete these prerequisites:

- Enable security console to work with permission groups
- Run scheduled processes to import security data
- (Optional) Assign privilege to use external REST API tools

Enable Security Console to Work with Permission Groups

For the Security Console to work with permission groups and related objects, set the **Enable Security Console External Application Integration** (ORA_ASE_SAS_INTEGRATION_ENABLED) profile option at the site level.

1. In the Setup and Maintenance work area, search for the **Manage Administrator Profile Values** task using the search link in the panel.
2. Search for the profile option and set the value for the **Site** profile level to **Yes**.

Run Scheduled Processes to Import Security Data

To import resources from LDAP, and transfer the necessary information into the security tables of Fusion Applications, run these two scheduled processes sequentially.

1. Import Resource Application Security Data
2. Import User and Role Application Security Data

You must run the process one after the other.

1. Go to **Navigator** > **Tools** > **Scheduled Processes**.
2. Click **Schedule New Process**.
3. Leave the type as **Job**.
4. Search for and select the process.
5. Submit the process.

Assign Privilege to Use External REST API Tools

To create External REST API tools in AI Agent Studio, the Create and Edit Backends for Visual Builder Studio (ORA_FND_TRAP_PRIV) privilege must be added to the custom role assigned to the user. You can add this privilege while creating or editing a custom role.

1. Go to **Navigator** > **Tools** > **Security Console**.
2. To use a new custom role, create it. To use an existing custom role, search for the custom role and edit it.
3. Go to the Function Security Policies page and select **Add Function Security Policy**.
4. Add the Create and Edit Backends for Visual Builder Studio (ORA_FND_TRAP_PRIV) privilege to the role and save it.

5. Save the custom role and assign to the user.

Assign Predefined Duty Roles to Job Roles

Assign predefined product-specific duty roles to the appropriate job roles, and make sure permission groups are enabled. You can give people access to configure AI agents in all or specific products.

> **CAUTION:** Using predefined roles might account for subscription consumption irrespective of whether you purchased the cloud service or not. See *Guidance for Assigning Predefined Roles*.

- *Provide Access to Configure AI Agents in all Products*
- *Provide Access to Configure AI Agents in Oracle Fusion Cloud Human Capital Management*
- *Provide Access to Configure AI Agents in Oracle Fusion Cloud Supply Chain & Manufacturing*
- *Provide Access to Configure AI Agents in Oracle Fusion Cloud Procurement*
- *Provide Access to Configure AI Agents in Oracle Permitting and Licensing*

Provide Access to Configure AI Agents in all Products

1. Go to **Navigator** > **Tools** > **Security Console** and create a new custom job role.

   > **Note:** Make sure to enable permission groups.

2. On the Role Hierarchy page, open the **Roles and Permission Groups** tab and add these duty roles:
   - Fai Genai Agent CX Administrator Duty (ORA_DR_FAI_GENERATIVE_AI_AGENT_CX_ADMINISTRATOR_DUTY)
   - Fai Genai Agent FIN Administrator Duty (ORA_DR_FAI_GENERATIVE_AI_AGENT_FIN_ADMINISTRATOR_DUTY)
   - Fai Genai Agent GRC Administrator Duty (ORA_DR_FAI_GENERATIVE_AI_AGENT_GRC_ADMINISTRATOR_DUTY)
   - Fai Genai Agent HCM Administrator Duty (ORA_DR_FAI_GENERATIVE_AI_AGENT_HCM_ADMINISTRATOR_DUTY)
   - Fai Genai Agent PRC Administrator Duty (ORA_DR_FAI_GENERATIVE_AI_AGENT_PRC_ADMINISTRATOR_DUTY)
   - Fai Genai Agent PRJ Administrator Duty (ORA_DR_FAI_GENERATIVE_AI_AGENT_PRJ_ADMINISTRATOR_DUTY)
   - Fai Genai Agent PSC Administrator Duty (ORA_DR_FAI_GENERATIVE_AI_AGENT_PSC_ADMINISTRATOR_DUTY)
   - Fai Genai Agent SCM Administrator Duty (ORA_DR_FAI_GENERATIVE_AI_AGENT_SCM_ADMINISTRATOR_DUTY)
3. If your user doesn't have the Application Implementation Consultant (ORA_ASM_APPLICATION_IMPLEMENTATION_CONSULTANT_JOB) job role, open the **Roles and Privileges** tab and add the Manage All Intelligent Agents (ORA_FAI_MANAGE_ALL_AI_AGENTS) role.
4. Save the custom role and assign it to users who want access.

Provide Access to Configure AI Agents in Oracle Fusion Cloud Human Capital Management

To give access to users without the Human Capital Management Application Administrator Job Role:

1. Go to **Navigator** > **Tools** > **Security Console** and create a new custom job role.

   > **Note:** Make sure to enable permission groups.

2. Go to the Role Hierarchy page.
   - Open the **Roles and Privileges** tab, and add the Manage HCM Intelligent Agent (ORA_HRC_HCM_AI_AGENT_MANAGEMENT_DUTY) duty role.
   - Open the **Roles and Permission Groups** tab, and add the Fai Genai Agent HCM Administrator Duty (ORA_DR_FAI_GENERATIVE_AI_AGENT_HCM_ADMINISTRATOR_DUTY) role.
3. Save the custom role and assign to the appropriate job roles.

To give access to users with the Human Capital Management Application Administrator Job Role:

1. Go to **Navigator** > **Tools** > **Security Console** and create a new custom job role.

   > **Note:** Make sure to enable permission groups.

2. On the Role Hierarchy page, open the **Roles and Permission Groups** tab, and add the Fai Genai Agent HCM Administrator Duty (ORA_DR_FAI_GENERATIVE_AI_AGENT_HCM_ADMINISTRATOR_DUTY) role.
3. Save the custom role and assign to the appropriate job roles.

Provide Access to Configure AI Agents in Oracle Fusion Cloud Supply Chain & Manufacturing

To give access to users without the Supply Chain Application Administrator Job Role:

1. Go to **Navigator** > **Tools** > **Security Console** and create a new custom job role.

   > **Note:** Make sure to enable permission groups.

2. On the Role Hierarchy page, open the **Roles and Privileges** tab and add these roles:

- SCM Intelligent Agent Management Duty (ORA_RCS_SCM_AI_AGENT_MANAGEMENT_DUTY)
- SCM Intelligent Agent Management Duty (ORA_RCS_SCM_AI_AGENT_MANAGEMENT_DUTY_HCM)

3. Open the **Roles and Permission Groups** tab and add the Fai Genai Agent SCM Administrator Duty (ORA_DR_FAI_GENERATIVE_AI_AGENT_SCM_ADMINISTRATOR_DUTY) duty role.
4. Save the custom role and assign this custom role to the appropriate job roles.

To give access to users with the Supply Chain Application Administrator Job Role:

1. Go to **Navigator** > **Tools** > **Security Console**.
2. Search for the Supply Chain Application Administrator (ORA_RCS_SUPPLY_CHAIN_APPLICATION_ADMINISTRATOR_JOB) job role, and edit it.
3. Enable permission groups and save the job role.

Provide Access to Configure AI Agents in Oracle Fusion Cloud Procurement

To give access to users without the Procurement Application Administrator Job Role:

1. Go to **Navigator** > **Tools** > **Security Console** and create a new custom job role.

   > **Note:** Make sure to enable permission groups.

2. On the Role Hierarchy page, open the **Roles and Privileges** tab and add these roles:
   - PRC Intelligent Agent Management Duty (ORA_PO_PRC_AI_AGENT_MANAGEMENT_DUTY)
   - PRC Intelligent Agent Management Duty (ORA_PO_PRC_AI_AGENT_MANAGEMENT_DUTY_HCM)
3. Open the **Roles and Permission Groups** tab and add the Fai Genai Agent PRC Administrator Duty (ORA_DR_FAI_GENERATIVE_AI_AGENT_PRC_ADMINISTRATOR_DUTY) duty role.
4. Save the custom role and assign this custom role to the appropriate job roles.

To give access to users with the Procurement Application Administrator Job Role:

1. Go to **Navigator** > **Tools** > **Security Console**.
2. Search for the Procurement Application Administrator (ORA_PO_PROCUREMENT_APPLICATION_ADMIN_JOB) job role, and edit it.
3. Enable permission groups and save the job role.

Provide Access to Configure AI Agents in Oracle Permitting and Licensing

To give access to users without the PSC Application Administrator Job Role:

1. Go to **Navigator** > **Tools** > **Security Console** and create a new custom job role.

   Note: Make sure to enable permission groups.

2. On the Role Hierarchy page, open the **Roles and Privileges** tab and add these roles:
   - PSC Intelligent Agent Management (ORA_PSC_AI_AGENT_MANAGEMENT_DUTY)
   - PSC Intelligent Agent Management (ORA_PSC_AI_AGENT_MANAGEMENT_DUTY_HCM)
3. Open the **Roles and Permission Groups** tab and add the Fai Genai Agent PSC Administrator Duty (ORA_DR_FAI_GENERATIVE_AI_AGENT_PSC_ADMINISTRATOR_DUTY) duty role.
4. Save the custom role and assign this custom role to the appropriate job roles.

To give access to users with the PSC Application Administrator Job Role:

1. Go to **Navigator** > **Tools** > **Security Console** and create a new custom job role.

   Note: Make sure to enable permission groups.

2. On the Role Hierarchy page, open the **Roles and Permission Groups** tab and add the Fai Genai Agent PSC Administrator Duty (ORA_DR_FAI_GENERATIVE_AI_AGENT_PSC_ADMINISTRATOR_DUTY) duty role.
3. Save the custom role and assign this custom role to the appropriate job roles.

3 Before You Begin

Choose How to Create AI Agents

Here are some examples of how you choose between creating agents from scratch or using preconfigured templates, based on your use case.

Use Case	How to Create AI Agents
You want to turn on Benefits Agent, which allows employees to ask questions about their benefits, within AI Agent Studio. To ensure that this agent can understand the benefits specific to the organization, the documents specific to that organization must be uploaded.	Create an AI agent using a preconfigured agent team template. You can edit the existing agent, without adding any additional tools. **Note:** You can either use or copy a template and edit it. When you copy an agent template, you can automatically add a suffix to all the components to easily differentiate them.
You want to create a custom agent to help with onboarding by answering new hire questions and providing deep links to internal new hire resources if the agent is unable to answer. You can tailor the agent's topics and prompts to ensure they effectively address the specific needs of new hires.	Create a custom agent from scratch, and add the necessary tools and topics.
You want to create a Contract Assistant Agent to generate a new contract with auto-filled fields. This assistant will be provided with multiple contract templates to generate the new contract. The agent also reviews existing contracts to perform risk analysis, recommend changes and route the contract for final approvals. For these activities that need distinct expertise (template management, legal validation, recommendation, and routing), the administrator will need to build a supervisor agent that collaborates with other worker agents.	You can use either of these methods: • Create a team of custom agents from scratch, with one agent added as a supervisor to manage the others. • Use a preconfigured agent team template, and add more agents and a supervisor agent to it.
You want to create a Supplier Quote Assistant agent to parse through quotes from multiple suppliers. The agent processes the quote document and maps the attributes. It then creates a draft	Create an agent team of type Workflow, to run the predetermined set of tasks. Each node can perform a defined function, for example, extracting data, calling a business object, running an LLM, or sending an email in the order you define.

Use Case	How to Create AI Agents
requisition and submits for approval. For these activities, administrator will need to create a workflow agent that runs the tasks in the specified order.	

Get Started with AI Agent Studio

You can start with a preconfigured agent template or create your own agent team.

Here's a broad outline of the tasks involved in creating an agent team. You can edit these same components when using a preconfigured agent template.

Task	Details
Define tools	To effectively define the tools required by an agent, you need to first identify the types of questions users might ask and then decide which tools the agent needs to answer those questions accurately. These are the available tools: • Oracle Fusion Cloud Applications business object - These business objects are data fields stored within Fusion Applications that AI agents can access as a part of agent configuration. Examples include the cost of a specific item, the time period that the benefits cover, and the date when a specific field was last updated by the user. AI agents can access Fusion Applications business objects, APIs, and rules without more configurations. You can control what data the agents can access. Just pick the business objects and their fields to use or ignore. AI agents respect native Fusion Applications security and role-based access controls, ensuring enterprise data protection and privacy. • Document upload - You can upload specific documents to be used by the AI agent, and provide natural language instructions on how the agent should use these documents. The agent can then search for information in the documents to provide a more exact answer to a user's question. • User profile - You can get information about the logged-in user, such as their geographic region or tenure at the company. • Calculator - The calculator performs numeric computations, such as calculating time-off balances. • Email - This tool can access the email client to send emails that include summaries of interactions or details pulled from a knowledge store. • Deep link - A deep link will send a user directly to the part of Fusion Applications where they can update underlying information. For example, if the user moved and wants to update their home address in the HR system, a deep link can quickly route the user to the page for making that update.
Define topics	Topics define the focus of the agent to a specific area of expertise. For example, within Benefits Administrator agent, we might define topics such as health policy coverage, vision policy coverage, and benefits enrollment. Use Topics to efficiently streamline your interactions with the agent.

Task	Details
	• Specify the instructions that help the agent decide which tools to use. • Enable the agent to better understand user intent by letting it identify and select the relevant topic based on the user's question. • Give each topic a clear, specific name, and include natural language instructions to ensure it's used correctly. You can reuse topics across agents.
Build new agents	Define the capabilities and scope of the agent. • Agent name • Product area the agent will work in • Natural language instructions to allow the agentic flow, or other agents, to understand the capabilities of this particular agent
Add a user (human) in the loop	If required, add an approval step for some actions that your AI agent will perform. The review step can be added at any point in the process for oversight and control over key actions, such as sending an email or updating a record.
Build the agent team	Create an agent team and add agents and other artifacts to it. Types of agent teams are: • Supervisor - A supervisor agent manages other agents and artifacts in the agentic flow. • Workflow - An agentic flow that does the tasks in a predetermined order. Agents and artifacts are added as nodes in the workflow, and each node performs a defined function, for example, extracting data, calling a business object, running an LLM, or sending an email. The node then passes its output to the next step.
Test the agents	Make sure to test the agent before deploying to production. Ask a test query, and determine the accuracy and relevance of the agent's response. You can also see the instructions the agent is following, and the actions the agent has taken to arrive at the response.
Deploy the agent team	After defining and testing your agent team, you can deploy it directly from AI Agent Studio. • Embed the agent conversation chat experience into any website or application. • Trigger the agent from an external resource using Webhooks, or seamlessly embed the chat experience into HTML and React web pages.

RAG Agents to AI Agent Studio

If you've previously created any RAG agents in Fusion Applications, we recommend replacing your existing agent with a new one created in AI Agent Studio. For details, see *Migrate Document Tools of RAG Agents*.

4 Create AI Agents

Create AI Agents Using Preconfigured Agent Team Templates

1. Go to **Navigator** > **Tools** > **AI Agent Studio**.
2. Select **Use Template** from the required agent team and provide details for the new agent team.

> **Tip:** To automatically add a suffix to all artifacts in your agent team, choose **Copy Template** instead of **Use Template**. When you copy a template, you're taken directly to the agent team canvas where you can edit the agent team settings, agents, tools, and topics. The **Use Template** option takes you through a step-by-step process for configuring each artifact in the agent team.

Details Tab

Field	Description
Family	Select the family to which this agent team belongs.
Product	Select the product within the family to which this agent team belongs.

Field	Description
Maximum Interactions	Indicate the number of times an agent within this agent team can interact with the topics and tools assigned to it.

LLM Tab

Field	Description
Provider	Select the LLM for your agent.

Security Tab

Field	Description
Add	Select the roles which will have access to this agent team.

Questions Tab

Field	Description
Starter Questions	Enter initial questions for the agent team.
Follow-up Questions	Enable this to indicate that the agent team can ask follow-up questions based on the user's conversation history.
Prompt	Enter the prompt to be used for the follow-up questions. For example: **using {chatHistory} generate 3 follow up questions in json format. JSON Schema format mentioned below. Remove the ``json markdown from the output. Here is the JSON Schema format the output should adhere to:** **[{"question": "<put first question generated here>"}, {"question": "<put second question generated here>"}, {"question": "<put third question generated here>"}]**
Insert Expression	Add additional variables to the prompt. For example, to add the current system date to your prompt, select the **Current Date Time** option.

Output Tab

Define the overall structure of the agent's output using JSON schema, to specify the exact output.

Field	Description
Specification Mode	Select this mode to directly modify the JSON schema for the output.

Field	Description
Simple Mode	Select this mode to define the output values and types. The corresponding JSON schema will be generated automatically and displayed in the specification mode for any further changes.

3. Continue to edit and add details for the agent and other artifacts such as tools and business objects in the agent team.

   > **Note:** The agents and other artifacts included within an agent template are optimized to provide the best usage of the agent team. We recommend not to change the basic functionality of the artifacts because that might impact your agent team's performance.

4. If needed, use ▷ to test the agent team. For any required fine-tuning, you can edit the agent team using ⚙.
5. Publish your agent team.

Users can view the published agentic flows from the AI Agents page. To open this page, add **agent-explore** to the end of the URL for AI Agent Studio. For example, **https://example.com/myApp/redwood/human-resources/ai-studio/agent-explore**.

Create Custom AI Agents of Type Supervisor from Scratch

1. Go to **Navigator** > **Tools** > **AI Agent Studio**.
2. From the Tools tab, add the required tools.

   For example, to create an HR benefits administrator agent that can answer questions related to medical, vision, dental, retirement, and stock plans, these are some of the tools needed.
   - Document tools, with the organization's health and financial benefits documents.

     > **Note:** Before you add any document tools to an agent, make sure you've done these things:
     > - Set the status of the document tool as **Ready to Publish**.
     > - Run the **Process Agent Documents** scheduled process from the Scheduled Processes work area.
     > - After the scheduled process completes, set the status of the document tool as **Published**.

   - Business object tools, to fetch the employee enrollments data.
   - Calculator tool, to check balance amounts and percentages.
3. From the Topics tab, add the required topics. Make sure you include instructions about these key areas.
   - What the topic is about
   - Tools to use with the topic
   - Examples of possible questions
   - Any guidelines and guardrails

4. From the Agents tab, add an agent.
    - In **Maximum Interactions** field, specify the number of times the agent can interact with the topics and tools assigned to it.
    - Describe the persona and role of the agent, including the tone to use.
    - Add the prompt for the agent, and select any variables from **Insert Expression**.
5. Assign the required tools and topics to the agent, and create it.
6. From the **Agent Teams** tab, add an agent team and provide details for the new agent team.

    - **Details Tab**

Field	Description
Family	Select the family to which this agent team belongs.
Product	Select the product within the family to which this agent team belongs.
Type	Select **Supervisor** to create an agent team where one agent acts as supervisor and manages other agents and artifacts.
Maximum Interactions	Indicate the number of times an agent within this agent team can interact with the topics and tools assigned to it.

    - **LLM Tab**

Field	Description
Provider	Select the LLM for your agent team.

    - **Security Tab**

Field	Description
Add	Select the roles which will have access to this agent team.

    - **Questions Tab**

Field	Description
Starter Questions	Enter initial questions for the agent team.
Follow-up Questions	Enable this to indicate that the agent team can ask follow-up questions based on the user's conversation history.
Prompt	Enter the prompt to be used for the follow-up questions. For example: using {chatHistory} generate 3 follow up questions in json format. JSON Schema format mentioned below. Remove the ``json markdown from the output. Here's the JSON Schema format the output should adhere to:

Field	Description
	[{"question": "<put first question generated here>"}, {"question": "<put second question generated here>"}, {"question": "<put third question generated here>"}]
Insert Expression	Add additional variables to the prompt. For example, to add the current system date to your prompt, select the Current Date Time option.

- **Output Tab**

Field	Description
Specification Mode	Select this mode to directly modify the JSON schema for the output.
Simple Mode	Select this mode to define the output values and types. The corresponding JSON schema will be generated automatically and displayed in the specification mode for any further changes.

7. Select **Create**.
8. Select and click **New Supervisor Agent**.
9. Enter details for the supervisor agent and create it. This supervisor agent is specific only to the agent team in which it's created, and can't be reused.
10. Select and add any existing agents as worker agents, or create worker agents.

    When you create worker agents from within the agent team, those agents are available only for this agent team.
11. Add any needed artifacts such as tools and topics.
12. If needed, use to test the agent team. For any required fine-tuning, you can edit the agent team using .
13. Publish your agent team.

Users can view the published agents from the AI Agents page. To open this page, add **agent-explore** to the end of the URL for AI Agent Studio. For example, **https://example.com/myApp/redwood/human-resources/ai-studio/agent-explore**.

Create Custom AI Agents of Type Workflow

A Workflow is a sequence of nodes agent teams of type Workflow use for deterministic, rule-based orchestration of tasks in which every step is preconfigured. Workflows are ideal for scenarios where compliance, repeatability, and governance are essential. The workflow tasks are represented as a connected sequence of nodes. Each node within the workflow performs a defined function, such as extracting data, calling a business object function, running an LLM, or sending an email. The output from each node is seamlessly passed to the subsequent step, ensuring a controlled and efficient process flow.

Let's look at an example to illustrate how a workflow operates in practice. Imagine a scenario where files are uploaded and need to be processed based on their current status:

- If a file is currently being indexed, the workflow extracts the file's content and stores it in a vector database.

- If a file has already been indexed, the workflow retrieves relevant content directly from the vector database.

In both cases, the final response to the user is generated by an LLM. This image shows how the workflow nodes are arranged for this use case.

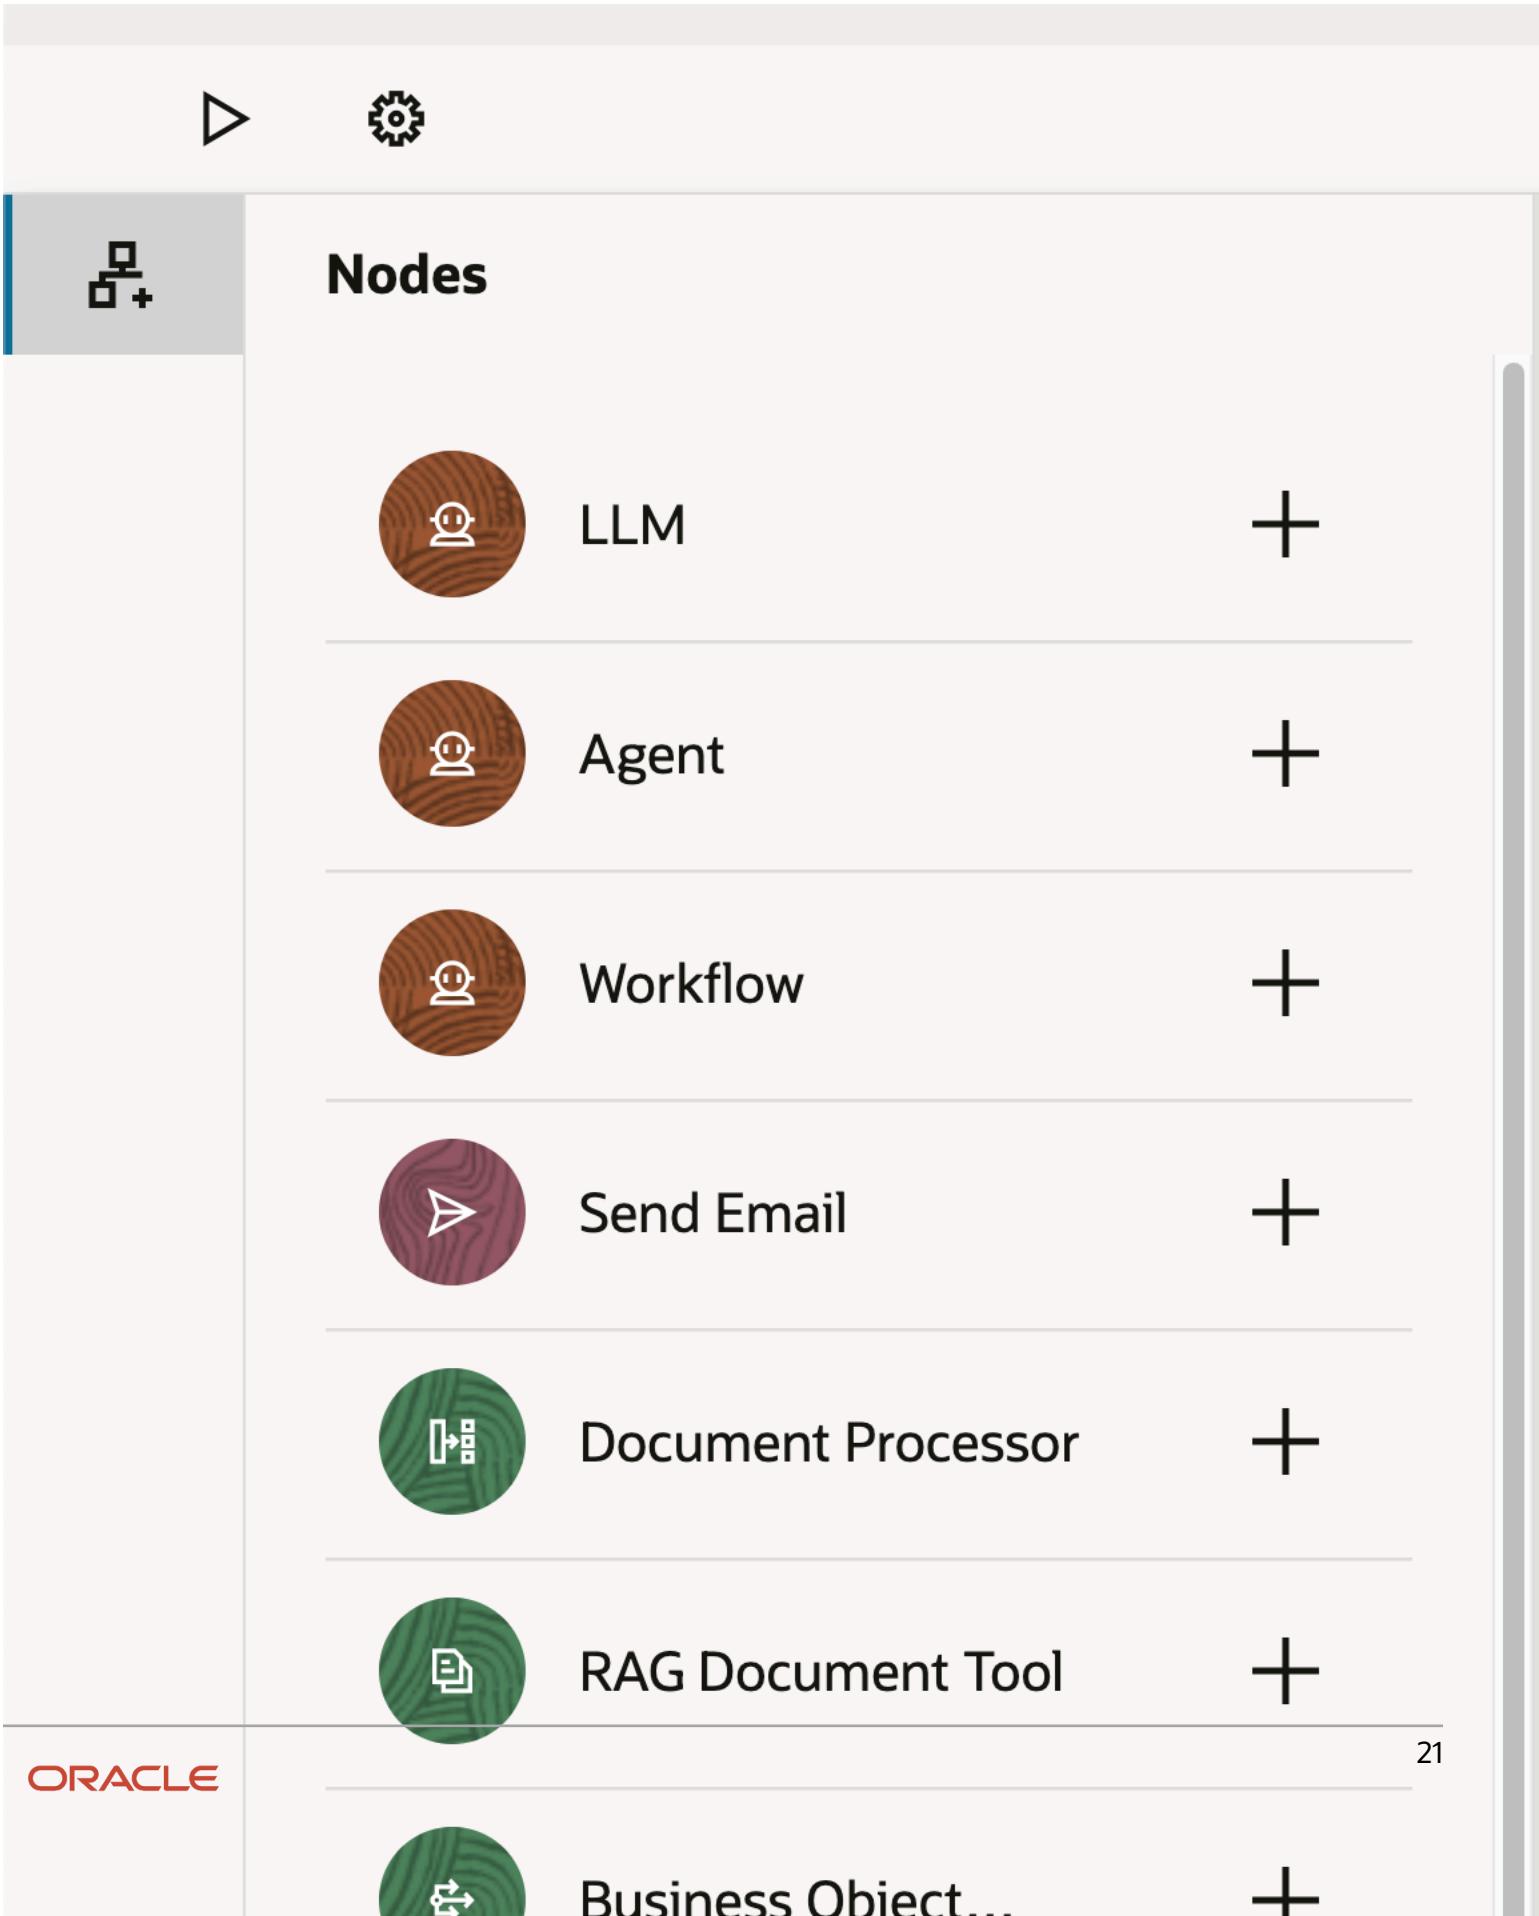
Nodes
LLM
Agent
Workflow
Send Email
Document Processor
RAG Document Tool
Business Object...

Callout Number	Node Name	Node Type	Description
1	defaults	Set Variables	Sets the default values for the variables specified, such as *attachmentId*, *document* uploaded document in the WebCenter Content server.
2	indexing	If Condition	Decides whether or not to index the information in the document.
3	file node	Document Processor	If *indexing* is **true**, processes the file by extracting its text, which is then stored in a
4	VECTOR_WRITER	Vector DB Writer	Stores the extracted file contents from the previous node in the vector database.
5	VECTOR_READER	Vector DB Reader	If *indexing* is **false**, retrieves the content directly by performing a semantic search database.
6	golden answer	LLM	Uses the LLM to process information extracted from the vector reader, generate ar the desired output.

Here are the steps to create your own workflow.

1. Go to **Navigator** > **Tools** > **AI Agent Studio**.
2. From the **Agent Teams** tab, add an agent team and provide details for the new agent team.

   Details tab

Field	Description
Family	Select the family to which this agent team belongs.
Product	Select the product within the family to which this agent team belongs.

Field	Description
Type	Select **Workflow**.

LLM tab

Field	Description
Provider	Select the LLM for your agent team.

Security tab

Field	Description
Add	Select the roles which will have access to this agent team.

Triggers tab

Create triggers to initiate the workflow. Triggers define the conditions or events that start the workflow and determine when and why it begins.

Field	Description
Type	Select the data type of the trigger.

Variables tab

Define workflow-level variables, making them accessible to all nodes within the workflow. Variables are ideal for storing IDs, constants, or any values you need to share and use across different parts of the workflow.

Field	Description
Type	Select the data type of the variable.

Output tab

Define the overall structure of the workflow's output using JSON schema, to specify exactly what you will receive when the workflow is completed. If no schema is specified, the output type defaults to that of the individual node. When a schema is provided, the system validates the outputs against the schema and enables auto-complete for those fields.

Field	Description
Specification Mode	Select this mode to directly modify the JSON schema for the output.

Field	Description
Simple Mode	Use this mode to define output values and types. The corresponding JSON schema will be generated automatically and displayed in the specification mode for you to edit.

Error Handling tab

Specify where to send an email when a workflow instance encounters a permanent error. You can also use context expressions in these fields for additional clarity.

3. Click **Create**.
4. Add nodes to the workflow by hovering over the line between the start and end of the workflow. Click to choose the type of node you want to add.

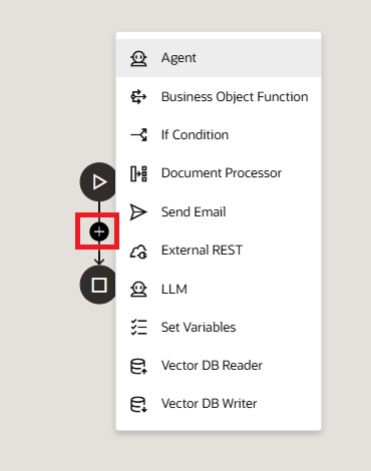

Tips

- Connect nodes in sequence to control how inputs are received, processed, and forwarded throughout the workflow.
- You can insert new nodes before, after, or between existing nodes as needed.
- Once created, you can't rearrange the nodes. To change a node's position, delete it and create it again in the desired spot.

5. If needed, use to test the agent team. For any required fine-tuning, you can edit the agent team using .
6. Publish the agent team.

Users can view the published agentic workflows from the AI Agents page. To open this page, add agent-explore to the end of the URL for AI Agent Studio. For example, `https://example.com/myApp/redwood/human-resources/ai-studio/agent-explore`.

5 Monitor and Evaluate AI Agents

Monitor and Evaluate AI Agents

Monitor and gain insights into how your AI agents are performing, and also evaluate the agents for accuracy. You can also track the interactions with your agents, understand real-world usage patterns, identify common errors, and measure overall performance.

- Monitoring - Monitoring tracks performance and provides insight on how your agents behave in production. Monitor agents to ensure that your quality bars for response time and token counts are maintained over time. You can also see any errors logged here.
- Evaluation - Evaluate agents before you deploy them, to ensure that they're ready for production. Test your agents for response correctness, response time, and token usage to meet your quality standards. After making any changes to your agent, or after a model update, rerun evaluations to confirm that your agent continues to perform as expected. This proactive approach helps you maintain high-quality experiences for your users.

This table summarizes some key metrics, their descriptions, and their availability for monitoring or evaluation.

Metric	Description	Available to Evaluate	Available to Monitor
Error Rate	Percentage of user sessions that ended in an error.	Yes	Yes
Error Count	Total number of errors recorded.	Yes	Yes
Session Count	Total number of conversations initiated with agents.	Yes	Yes
P99 Latency	The maximum wait time in milliseconds for 99% of users, revealing any areas where you should review and optimize the prompts or structure of the agent.	Yes	Yes
P50 Latency	The maximum wait time in milliseconds for 50% of users, helping identify performance issues. You can view this metric in the details of the monitoring or evaluation results.	Yes	Yes
Total Tokens	Cumulative number of tokens used by all agents.	Yes	Yes

Metric	Description	Available to Evaluate	Available to Monitor
Input Token Count	Total tokens sent to the LLM for requests. This includes system prompts, user messages, retrieved or context data, chat history, and tool or function definitions.	No	Yes
Output Token Count	Total tokens generated by the LLM for requests sent to it.	Yes	Yes
Median Correctness	The 50th percentile of correctness scores across evaluation runs. Each score (0–1) is computed by comparing the agent's answer to the reference answer provided in the evaluation set.	Yes	No
Session Count	The number of unique conversational sessions between a user and an AI agent. One session can include multiple messages or evaluation runs.	Yes	Yes

Prerequisites

Aggregate the metrics that are displayed on the Monitoring and Evaluation tab in AI Agent Studio.

1. Go to **Navigator** > **Tools** > **Scheduled Processes**.
2. In Scheduled Processes, click **Schedule New Process**.
3. Leave the type as **Job**.
4. Search for and select **Aggregate AI Agent Usage and Metrics**.
5. Run the **Aggregate AI Agent Usage and Metrics** scheduled process.

   You can schedule this process to run on a recurring basis, for example, once a day.

The process aggregates the metrics that are displayed in the Monitoring and Evaluation tab of AI Agent Studio.

Monitor Agents

1. Go to **Navigator** > **Tools** > **AI Agent Studio**.
2. From the Monitoring and Evaluation tab, go to the Monitoring tab.
   - The Monitoring subtab displays aggregated metrics of all the agent runs over the selected time frame.
   - All agent runs are monitored, including agents in draft status.

3. Select the agent to view additional details.

   Each row represents a single session and displays the number of turns (back-and-forth messages), the session's completion status (successful or error), and the number of tokens used.
4. Select any session in the list to open the detailed trace view. This view displays a step-by-step timeline of the entire conversation, showing exactly which tools were called, the duration of each step, and the metrics for each step.

Evaluate Agents

Use evaluation sets to assess your agents' performance. An evaluation set contains one or more test questions, the expected agent responses, and the metrics to be measured. Evaluation sets are specific to each agent, and an agent can have multiple evaluation sets.

1. Create an evaluation set for an agent.
   a. Go to **Navigator** > **Tools** > **AI Agent Studio**.
   b. Select the Monitoring and Evaluation tab.

      All evaluations run on the agents are displayed in the Evaluation tab.
   c. To create an evaluation set, click **Manage Evaluations** and select  .
   d. Enter a name, code, and description for the evaluation set, and select the agent team to be evaluated.
   e. Choose the run mode.

      Sequential: Runs the questions in the exact order you define them. Use this if one question depends on the context of the previous one.

      Random: Runs the questions in a random order.
   f. From the Questions tab, add common questions that users are likely to ask the agent and the answers you would like the agent to deliver. Ensure both the questions and answers are concise, user-friendly, and reflect best practices.

      You can either add your questions and the expected answers individually, or upload a CSV file with the questions in the first column and the expected answers in the second column.
   g. From the **Metrics** tab, edit each metric to set the pass and fail criteria. For example, to indicate that the test fails if the correctness score is less than 0.7, choose **<** as the threshold condition and enter 0.7 as the threshold value.
   h. Select **Create** to save the evaluation set.
2. Run the evaluation set.
   a. On the Manage Evaluations page, select the **Initiate Evaluation Run** action for your evaluation set.
   b. Choose the version of the agent team to evaluate and run the evaluation.

3. Analyze the results.
    a. Click the evaluation set to view the Evaluation Runs page.
    b. Select the evaluation run and select the **View Run Results** action.

Tab	Information Displayed
Response performance	- Comparison of the expected response versus the actual response from the agent for each question, along with the metrics for each question in the evaluation. - Trace provides information about the detailed timeline for each question in the evaluation.
Correctness	Detailed breakdown of the correctness score. The LLM provides an initial score and feedback, and you can add your own feedback for record keeping in the **Correctness Score by Human** column.

Compare Evaluation Runs

You can see a side-by-side comparison of two different runs of the same evaluation, and easily spot regressions or improvements in latency, correctness, and token usage. Doing so, you can understand how an agent's performance changes over time, especially after you've made modifications.

1. From the Evaluation tab, select the evaluation.
2. Select any two runs and click **Compare**.
    - The Summary tab displays a high-level overview of the performance differences between the runs.
    - The Details tab provides a granular, question-by-question breakdown of the runs. For each question in the evaluation set, you can directly compare the actual response from Run 1 against the actual response from Run 2. You can also compare the specific latency, tokens used, and trace links for each question, making it easy to pinpoint exactly where and why performance or accuracy has changed.

6 Migrate RAG Agents to AI Agent Studio

Migrate Document Tools of RAG Agents

If you've created any RAG agents in Oracle Fusion Cloud Applications previously, we recommend that you replace your existing agent with an agent you create in AI Agent Studio.

You can migrate the Document tool you created for your existing agent to AI Agent Studio.

1. Go to the **Configure RAG Agents** work area.

   > **Tip:** You can use global search to get to the work area.

2. From the Tools tab, click  to migrate your tool.
3. Enter a unique tool name and code. That will help you easily find your migrated tool in AI Agent Studio.

After migrating your tool, create a new agent in AI Agent Studio using an appropriate template, add your migrated tool, and publish your agent. After testing and verifying your new agent, you can delete the original agent in the **Configure RAG Agents** work area.

7 AI Agent Studio FAQs

Why can't I see all agent teams in the Agent Teams tab of AI Agent Studio?

That's because only published agent teams are displayed by default. Use the **Draft** button to view agents that aren't published yet.

Why is my agent not fetching data from the document added to it?

Check if you've marked your document as published. If not, publish your document. To do so, from the Tools tab in AI Agent Studio, set the status of your document as **Ready to Publish**, and then run the **Process Agent Documents** scheduled process. After the scheduled process completes, from the Tools tab, select your document, and set its status as **Published**.

Also remember, the best practice is to create one Document tool per logical document. For example:

- Java Resumes Tool - One logical document containing all Java resumes.
- Python Resumes Tool - One logical document containing all Python resumes.

When should I use a single-agent flow vs a multiagent process?

- Use a single-agent process for straightforward business requirements, such as retrieving answers from a single policy document.
- Opt for a multiagent process when dealing with more complex scenarios. For example, when multiple documents need to be parsed, or when the answers vary based on the user asking the question. In these cases, multiple agents are used, and the user's question is directed to the appropriate worker agent for processing.

What types of files can I upload as part of Document tools in AI Agent Studio?

You can upload PDF, TXT, and Markdown files.

What's the difference between editing a topic versus editing a system prompt?

When you update a topic, the changes are available across agents because you can reuse topics across agents. When you're updating a system prompt, the changes are available only for the agent being edited.

Tip: We recommend updating topics instead of prompts, so that you can reuse the topics.

Why does the External REST Tool only support HTTPS transactions?

The tool is limited to only HTTPS to maintain data privacy in transit.

Why is the Roles list in the Security tab empty?

That's probably because permission groups aren't enabled for your duty role. Security administrators can enable permission groups for the duty role assigned to you, in Security Console. For details, see *Access Requirements for AI Agent Studio*.